

Pupil dilation indexes automatic and dynamic inference about the precision of stimulus distributions

Francesco Silvestrin*¹, Will D Penny^{1,2}, Thomas HB FitzGerald^{1,2,3}

1. *School of Psychology, University of East Anglia, Norwich Research Park, Norwich, Norfolk, NR4 7TJ, UK*
2. *The Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London, WC1N 3BG, UK*
3. *Max Planck-UCL Centre for Computational Psychiatry and Ageing Research, Russell Square House, London, WC1B 5EH, UK*

*Correspondence: Francesco Silvestrin

School of Psychology, University of East Anglia, Norwich Research Park, Norwich, Norfolk, NR4 7TJ, UK

E-mail: F.Silvestrin@uea.ac.uk

Telephone: +44(0)1603456161

Abstract

Learning about the statistics of one's environment is a fundamental requirement of adaptive behaviour. In this experiment we probe whether pupil dilation in response to brief auditory stimuli reflects automatic statistical learning about the underlying stimulus distributions. Specifically, we consider whether pupil dilation reflects automatic (task-irrelevant) learning about the precision of Gaussian distributions of pitch in a sequence of tones. We provide clear evidence, both by comparing responses to perceptually identical probe tones in low and high precision blocks, and using a novel model-based analysis, that subjects did indeed track the precision of the stimulus distribution. This extends previous work looking at electrophysiological effects of precision (or, equivalently, variance) learning, and provides new evidence that the putatively noradrenergic processes underlying pupil dilation reflect rapidly updated information about distributions of sensory stimuli. In addition, our study represents a validation of our model-based approach to analysing pupillometry data, which we believe has considerable promise for future studies.

Keywords: Auditory Oddball; Predictive Coding; Surprise; Statistical learning; Pupillometry; Precision

1. Introduction

Sensitivity to the statistics of one's environment is a key requirement for adaptive behaviour. Correspondingly, statistical learning has been a fertile area of research (e.g. Saffran et al., 1996, 1999; Kirkham et al., 2002), often in conjunction with probabilistic theories of cognition (Fiser, Berkes, Orbán, & Lengyel, 2010; Tenenbaum, Kemp, Griffiths, & Goodman, 2011; Turk-Browne, Scholl, Johnson, & Chun, 2010). Here, we consider statistical learning to solve the specific problem of forming beliefs about the precision (or, equivalently its inverse, the variance) of the underlying probability distributions that govern incoming sensory stimuli. From a normative perspective, accurately

estimating precision is extremely important, as it governs key features of learning and inference, such as how quickly to update one's beliefs (Behrens, Woolrich, Walton, & Rushworth, 2007; Mathys, Daunizeau, Friston, & Stephan, 2011) and how to weight incoming sensory information against prior beliefs (K. Friston, 2008), and aberrant precision estimation is widely believed to play a key role in psychopathology (Adams, Stephan, Brown, Frith, & Friston, 2013; Fletcher & Frith, 2008; K. J. Friston, Stephan, Montague, & Dolan, 2014; Lawson, Rees, & Friston, 2014)

To explore learning about precision, we make use of the widely replicated finding that non-luminance related pupil dilation indexes the surprise associated with incoming sensory stimuli (Alamia, VanRullen, Pasqualotto, Mouraux, & Zenon, 2019; Damsma & van Rijn, 2017; De Berker et al., 2016; Friedman, Hakerem, Sutton, & Fleiss, 1973; Kloosterman et al., 2015; Lavín, Martín, & Jubal, 2014; Nassar et al., 2012; O'Reilly et al., 2013; Preusschoff, 't Hart, & Einhäuser, 2011; Qiyuan, Richer, Wagoner, & Beatty, 1985; Raisig, Welke, Hagedorf, & van der Meer, 2010; Reinhard & Lachnit, 2002). (Here we define surprise as the negative log probability of an event occurring, though see (Baldi & Itti, 2010; Philipp Schwartenbeck, FitzGerald, & Dolan, 2016; Zenon, 2019) for an important alternative). This permits one to make inferences about participants' implicit beliefs about the statistics of their environment, without the necessity for an explicit probe or decision, and thus provides a means to characterise statistical learning processes (Alamia, VanRullen, Pasqualotto, Mouraux, & Zenon, 2019; Vincent, Parr, Benrimoh, & Friston, 2019). Specifically, where predictions are more precise, subjects should be more surprised by objectively identical stimuli, leading to a greater dilation response.

Learning about precision has previously been tested using reaction time and magnetoencephalography (MEG) data by (Garrido, Sahani, & Dolan, 2013). Here, subjects performed a modified version of the auditory oddball task, where on separate blocks tones were drawn from either high or low precision Gaussian distributions in log frequency space (Figure 1). The present work seeks to extend this approach, using pupillometry data. This is important, given the tight link between pupil dilation and noradrenergic activity in the locus coeruleus (Murphy, Robertson, Balsters, & O'Connell, 2011), since it provides the opportunity to better understand the function of the noradrenergic system (Dayan & Yu, 2006). Pupillometry data is also much cheaper and easier to acquire than neuroimaging data, and its use to characterise statistical learning in different groups (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015; Montague, Dolan, Friston, & Dayan, 2012) is thus attractive from a practical perspective.

We thus collected pupillometry data whilst subjects performed a slightly modified version of the paradigm used by Garrido et al. (2013). We used the data to perform two types of analysis. The first was closely modelled on that used in previous work (Garrido et al., 2013), and directly compared responses to probe tones during high and low precision blocks. Our main hypothesis was that a deviant sound (probe tones at 2000 Hz) would be more surprising, and therefore elicit a bigger pupil dilation, in high precision compared to low precision blocks, as the associated probability density would be lower (Figure 1). In the second we combined agent-based modelling (O'Doherty, Hampton, & Kim, 2007) with a novel convolution-based approach to analysing pupillometry data (Denison, Parker, & Carrasco, 2019; Zénon, 2017), which we believe has considerable promise for exploring automatic statistical learning in humans. We hypothesised that pupil dilation responses would reflect dynamic updating of beliefs about the precision of the stimulus distribution.

2. Methods

2.1. Participants

16 participants (12 females), aged 18 to 34 (mean = 21.1) took part in the study and they all gave informed consent.

2.2.Procedure

Participants were asked to look at a fixation cross in the centre of a computer screen while listening to a series of tones through headphones. Their task consisted only in pressing the space bar when the sound came only from one of the two headphones' speakers (i.e. when they heard it coming only from one side). Importantly, the pitch of the tones was entirely task irrelevant, deconfounding outlier and target (unilateral) tones, which is a concern with several previous studies (Hong, Walz, & Sajda, 2014; Liao, Yoneya, Kidani, Kashino, & Furukawa, 2016), as well as indexing automatic, rather than task-dependent statistical learning processes. Furthermore, the absence of a visual target avoided luminance changes that could alter pupil diameter.

The experiment was divided into 4 sessions, during each of which subjects were presented with 800 pure tones, each lasting 50 ms, with an interstimulus interval of one second. 4 blocks (2 high precision blocks and 2 low precision blocks) were present in each session, with 200 tones each and no breaks between blocks. The order of the blocks was counterbalanced and participants were not aware of the presence of different blocks within each session. Stimuli were selected to be similar to those used in a prior study (Garrido et al., 2013). Specifically, for each session the frequency of 688 out of 800 tones was sampled from a Gaussian distribution in log-frequency space, with mean $\mu = 500$ Hz and standard deviation $\sigma_h = 0.5$ octaves for high precision blocks and $\sigma_l = 1.5$ octaves for low precision blocks. (These values correspond to precisions of 4 and 0.4444 respectively). Out of the 112 remaining tones, 56 were standard probes (500 Hz, corresponding to the mean of the distribution) and 56 were deviant probes (2000 Hz, two octaves above the mean), which slightly distorted the probability distribution, adding two point-masses of 7% probability each. The number of unilateral, target tones varied across sessions (82, 80, 75, and 78 respectively). Both probes and targets were pseudo-randomly inserted in the stream, and targets were made to never coincide with a probe, or occur immediately after it. This was done because targets are very likely to elicit a strong, long lasting pupil response, which would confound the effect of surprise.

See Figure 2 for a graphical representation of the tone sequence in low and high precision blocks, and of how probe tones distorted the distributions.

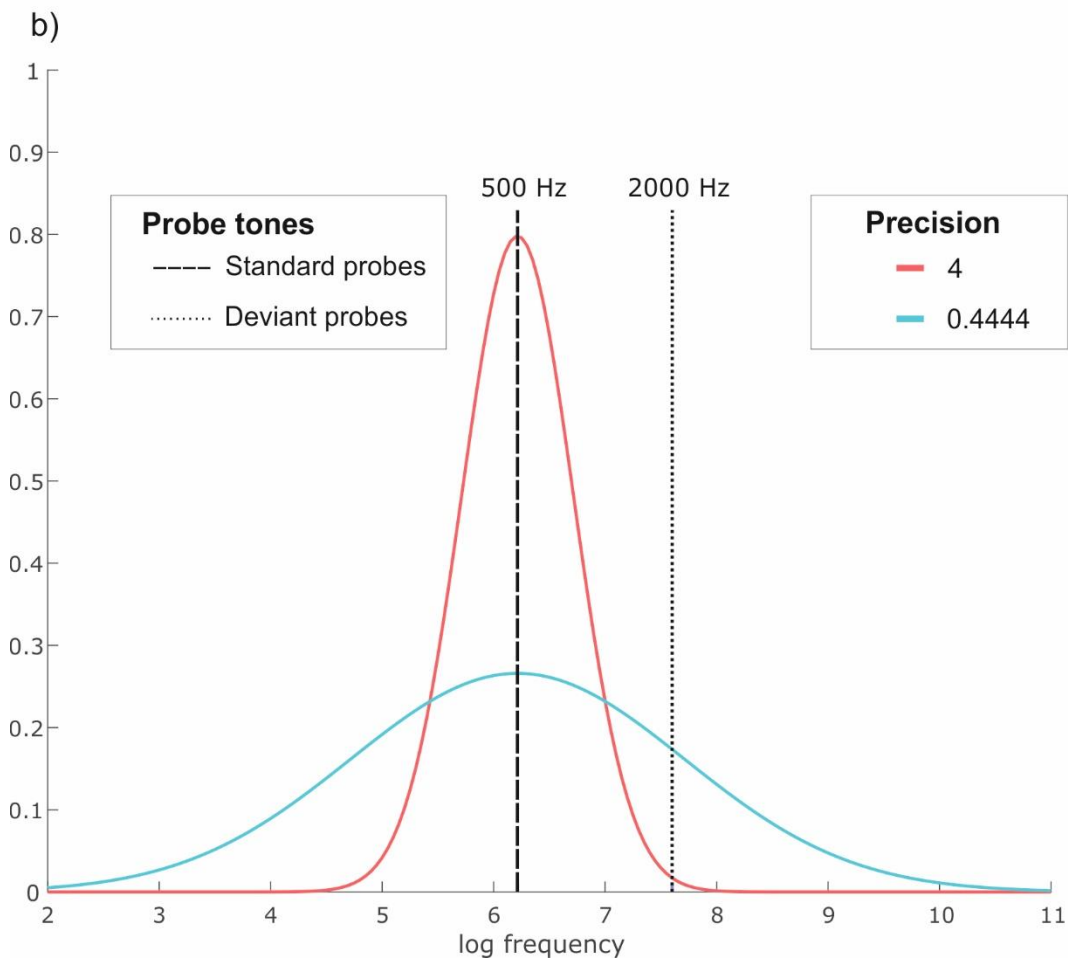
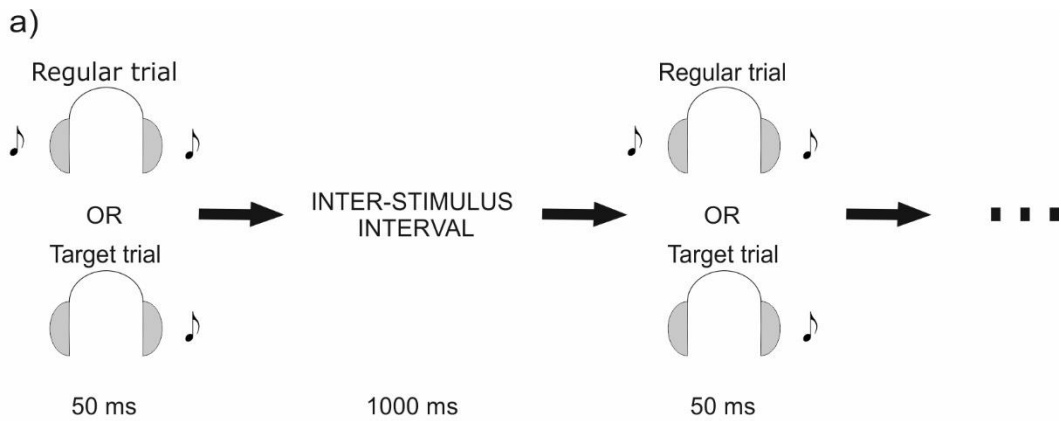


Figure 1: Illustration of experimental paradigm (a). Participants were exposed to a series of tones (800 per session, 3200 in total) and were asked to press the space bar when they heard the sound coming only from one speaker (i.e. only from one side). The pitch of the tones was sampled from two different probability distributions (b), alternating between high and low precision blocks. In line with previous work (Garrido et al., 2013) probes were added at 500 Hz and 2000 Hz, slightly distorting the distribution.

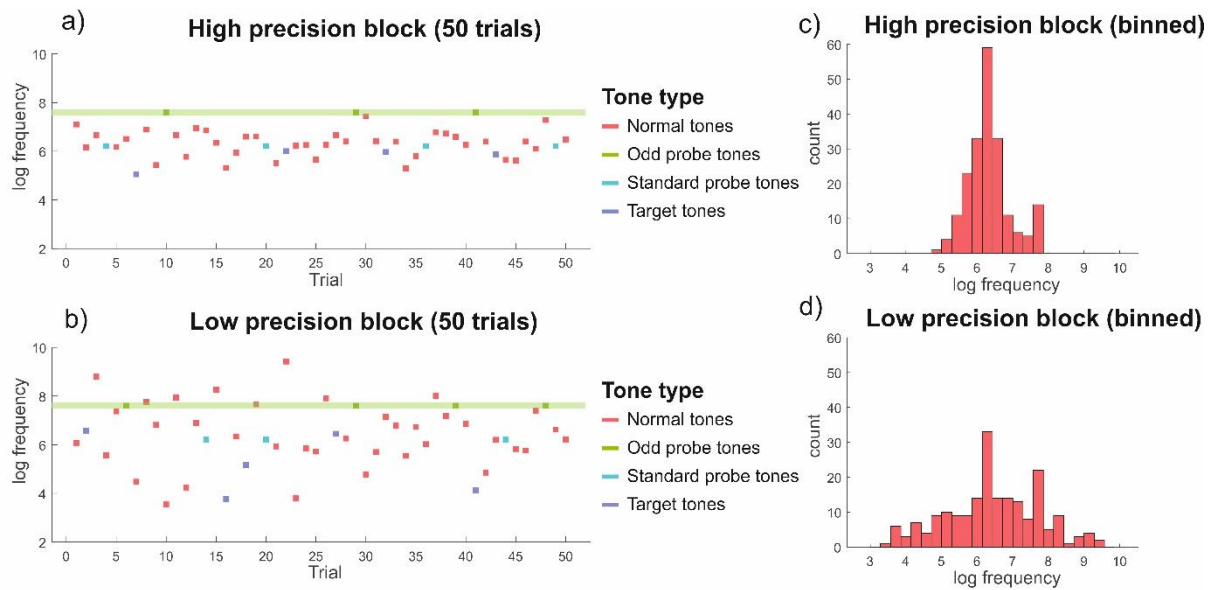


Figure 2: Graphical representation of 50 trials (corresponding to 50 tones) in a high (a) and low (b) precision block. Deviant probes (2000 Hz) are highlighted in green to evidence how they stand out more in high compared to low precision blocks. These and standard probes (500 Hz) slightly distorted the stimuli distribution, as shown in binned distributions (1 bin = 2/7 of an octave) of a sample high (c) and low (d) precision block (200 trials per block).

2.3. Pupillometry data recording and preprocessing

Pupillometry data were recorded at 500 Hz using an EyeLink 1000 eye-tracking device whilst subjects sat in a moderately lit room. As we were interested in learning effects rather than precise psychophysics, we did not directly measure the luminance of the screen or the fixation cross. Similarly, tones were presented at a constant volume, at a level that was comfortable but clearly audible for subjects, but these levels were not recorded. Critically, because these quantities were held constant for each subject throughout the duration of the experiment, purely physical properties of the stimuli or environment can not explain the results we describe below.

Linear interpolation was used to remove artefacts relating to eyeblinks and saccades, and the data were low pass filtered at 20 Hz. Data were recorded from both eyes simultaneously, and averaged prior to further analysis. Additionally, for the model-based analysis, time series were downsampled to 50 Hz, and were then mean-corrected and normalised to unit variance, to standardise responses across subjects.

2.4. Probe tone analysis

We first analysed our data using a classical model-free approach, where we averaged the responses to the probe tones in each precision condition for each subject. We then averaged all data-points from 900 ms to 1000 ms after stimulus onset and performed a two-way repeated measures ANOVA with precision (high vs low) and probe type (deviant vs standard) as within-subjects factors. The 900 ms-1000 ms time window was averaged to avoid multiple comparisons while trying to capture the peak of the pupil response, which typically occurs around 930 ms from stimulus onset (Hoeks & Levelt, 1993).

2.5. Time-series modelling

Building on previous work analysing fMRI time series (W. Penny, Kiebel, & Friston, 2003), we developed a model for analysing pupillometry data combining a General Linear Model (GLM) incorporating a convolution kernel, and an autoregression (AR) component, in order to model fluctuations not captured by the convolution model (which could come from a variety of sources) (Zénon, 2017). For simplicity, we only consider a first order AR model, but this approach can naturally be extended to include higher orders (W. Penny et al., 2003). (For related approaches, see Korn and Bach, 2016; Zénon, 2017; Vincent et al., 2019).

In this GLM-AR(1) model, data $\mathbf{z} = \{z_1, \dots, z_T\}$ are modelled in terms of a $[T \times K]$ design matrix \mathbf{X} , a $[K \times 1]$ vector of regression coefficients, and a $[T \times 1]$ vector of errors \mathbf{e} .

$$\begin{aligned} \mathbf{z} &= \mathbf{X}\mathbf{w} + \mathbf{e}, \\ e_t &= a(z_{t-1} - \mathbf{x}_{t-1} \cdot \mathbf{w}) + i_t, & \text{if } t > 1 \\ e_t &= r, & \text{if } t = 1 \end{aligned} \quad (1)$$

Here e_t is modelled as a combination of the ‘prediction error’ at the previous time point, weighted by the AR coefficient a , and i_t which is drawn from \mathbf{i} , a $[T \times 1]$ vector of Independent and Identically Distributed (IID) errors. The model also includes an additional session-specific parameter r to model the error at $t = 1$.

The design matrix \mathbf{X} is generated by convolving a $[T \times K]$ input matrix \mathbf{U} with a gamma kernel to model the slow time course of pupil responses (Denison et al., 2019; Hoeks & Levelt, 1993; Korn & Bach, 2016). The kernel is parameterised using three parameters: shape (h) and scale (l) parameters governing the properties of the Gamma distribution, and a delay parameter (d) introducing a temporal delay. Thus:

$$\begin{aligned} g &= \text{Gamma}(t - d; h, l), \\ \mathbf{x}_{\bullet k} &= \mathbf{u}_{\bullet k} * g. \end{aligned} \quad (2)$$

where Gamma indicates the probability density function of the gamma distribution. Our approach thus naturally accounts for between-subject variability in the time course of pupil responses, since the shape of the gamma distribution is fitted to individual responses (Denison et al., 2019) (Note that in our experiment there were no luminance changes, and we thus needed only to model pupil dilation responses. The GLM-AR approach is equally applicable to data involving pupil constriction (Korn & Bach, 2016), and we will consider this in future work).

2.6 Cognitive modelling

To model cognitive processes during the task, we devised a simple agent where at each trial i , trial-specific mean (m_i) and log-precision (k_i) estimates were generated using a predictive coding algorithm (Friston & Kiebel, 2009; Rao & Ballard, 1999; Shipp, 2016; Spratling, 2017), augmented to allow for dynamic estimation of precision. This does not imply any strong claim about the actual neuronal inference mechanisms employed by subjects on the task, since other schemes could undoubtedly provide similar predictions (Aitchison & Lengyel, 2016; K. Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017; Ma, Beck, Latham, & Pouget, 2006). It does, however, provide a simple and parsimonious method for modelling task performance, which is grounded in computational modelling of cortical function.

In this model, the agent infers on the log joint probability of current hidden states and observations, given all current and previous observations $\mathbf{y}_{1:i} = \{y_1, \dots, y_i\}$, prior beliefs about the initial hidden states (m_0 and k_0) and two fixed parameters which govern how quickly the mean ($\lambda^{(m)}$) and log precision ($\lambda^{(k)}$) are expected to change (often known as their volatility (Behrens et al., 2007; Mathys et al., 2011)). (Note that these parameters were fitted to individual subjects' data, meaning that we can capture a broad spectrum of beliefs about volatility).

$$\ln p(m_i, k_i | \mathbf{y}_{1:i}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) \quad (3)$$

Assuming the Markov property, this can be rewritten as

$$\ln \int \int p(m_i, k_i | y_i, m_{i-1}, k_{i-1}, \lambda^{(m)}, \lambda^{(k)}) p(m_{i-1}, k_{i-1} | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) dm_{i-1} dk_{i-1} \quad (4)$$

and estimated recursively, which represents standard Bayesian filtering.

The model has the conditional independence properties that

$$p(y_i | \mathbf{y}_{1:i-1}, \mathbf{m}_{1:i}, \mathbf{k}_{1:i}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) = p(y_i | m_i, k_i), \quad (5)$$

$$p(m_i, k_i | y_{1:i-1}, \mathbf{m}_{1:i-1}, \mathbf{k}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) = p(m_i | m_{i-1}, \lambda^{(m)}) p(k_i | k_{i-1}, \lambda^{(k)}) \quad (6)$$

meaning that observations depend only on the current hidden states, and each sequence of hidden states is a separate Markov chain.

More specifically, each continuous-valued observation y_i is sampled from a normal distribution with mean m_i and log-precision k_i

$$p(y_i | m_i, k_i) = N(y_i | m_i, e^{-k_i}), \quad (7)$$

where $N(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2 .

Similarly, the mean and log-precision of the distribution are treated as independent, zero-mean, Gaussian random walks, with fixed precisions (volatilities), given by

$$p(m_i | \mathbf{m}_{1:i-1}, \mathbf{k}_{1:i-1}, \lambda^{(m)}) = N(m_i | m_{i-1}, \lambda^{(m)-1}) \quad (8)$$

$$p(k_i | \mathbf{m}_{1:i-1}, \mathbf{k}_{1:i-1}, \lambda^{(k)}) = N(k_i | k_{i-1}, \lambda^{(k)-1}). \quad (9)$$

The log joint distribution can now be written as:

$$\begin{aligned} I &= \ln p(y_i, m_i, k_i | \mathbf{y}_{1:i}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) \\ &= \ln p(y_i | m_i, k_i) \dots \\ &\quad + \ln \left(\int \int p(m_i | m_{i-1}, \lambda^{(m)}) p(k_i | k_{i-1}, \lambda^{(k)}) p(m_{i-1}, k_{i-1} | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) dm_{i-1} dk_{i-1} \right), \end{aligned} \quad (10)$$

To avoid the need to estimate joint probability distributions, the agent performs variational inference (Beal, 2003), and approximates the log joint with a distribution $q(m_i, k_i)$ which factorises such that:

$$q(m_i, k_i) = q(m_i)q(k_i). \quad (11)$$

This gives a variational lower bound on the log model evidence

$$\begin{aligned} L &= \ln p(y_i | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) - D_{KL} \left(q(m_i, k_i) \parallel p(m_i, k_i | \mathbf{y}_{1:i}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) \right), \\ &= \iint q(m_i, k_i) \ln p(y_i | m_i, k_i) dk_i dm_i + \int q(m_i) \ln p(m_i | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) dm_i \dots \\ &\quad + \int q(k_i) \ln p(k_i | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) dk_i - \int q(m_i) \ln q(m_i) dm_i \\ &\quad - \int q(k_i) \ln q(k_i) dk_i. \end{aligned} \quad (12)$$

To generate a predictive distribution over m_i , the agent uses the variational posterior $q(m_{i-1})$ generated on the previous trial. Thus

$$\begin{aligned} p(m_i | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) &\approx \int p(m_i | m_{i-1}, \lambda^{(m)}) q(m_{i-1}) dm_{i-1} \\ &= N(m_i | \tilde{\mu}_i^{(m)}, \tilde{\tau}_i^{(m)-1}), \end{aligned} \quad (13)$$

$$\tilde{\mu}_i^{(m)} = \hat{\mu}_{i-1}^{(m)}, \quad (14)$$

$$\tilde{\tau}_i^{(m)} = \left(\lambda^{(m)-1} + \hat{\tau}_{i-1}^{(m)-1} \right)^{-1}, \quad (15)$$

where $\hat{\mu}_{i-1}^{(m)}$ and $\hat{\tau}_{i-1}^{(m)}$ are the sufficient statistics (mean and precision) of $q(m_{i-1})$, and $\tilde{\mu}_i$ and $\tilde{\tau}_i^{(m)}$ are the sufficient statistics of the predictive distribution. (At $i = 1$, $\tilde{\mu}_1 = m_0$ and $\tilde{\tau}_1^{(m)} = \lambda^{(m)}$).

Similarly, the predictive distribution over k_i used by the agent is given by

$$\begin{aligned} p(k_i | \mathbf{y}_{1:i-1}, m_0, k_0, \lambda^{(m)}, \lambda^{(k)}) &\approx \int p(k_i | k_{i-1}, \lambda^{(k)}) q(k_{i-1}) dk_{i-1} \\ &= N(k_i | \tilde{\mu}_i^{(k)}, \tilde{\tau}_i^{(k)-1}), \end{aligned} \quad (16)$$

$$\tilde{\mu}_i^{(k)} = \hat{\mu}_{i-1}^{(k)} \quad (17)$$

$$\tilde{\tau}_i^{(k)} = \left(\lambda^{(k)-1} + \hat{\tau}_{i-1}^{(k)-1} \right)^{-1} \quad (18)$$

(At $i = 1$, $\tilde{\mu}_1^{(k)} = k_0$ and $\tilde{\tau}_1^{(k)} = \lambda^{(k)}$).

Thus for the agent

$$L = \left\langle \ln N(y_i | m_i, e^{-k_i}) + \ln N(m_i | \tilde{\mu}_i^{(m)}, \tilde{\tau}_i^{(m)-1}) + \ln N(k_i | \tilde{\mu}_i^{(k)}, \tilde{\tau}_i^{(k)-1}) \right\rangle_{m_i, k_i} \dots \quad (19)$$

$$+ h(q(m_i)) + h(q(k_i)),$$

where $\langle \cdot \rangle$ denotes the expectation and $h(\cdot)$ the entropy of a probability distribution.

The optimal solution $q^*(\cdot)$ for variable m_i can now be derived simply using the properties of the Gaussian distribution and standard properties of the variational inference (see (Bishop, 2006) for a fuller exposition), and is given by

$$\begin{aligned} \ln q^*(m_i) &= \left\langle \ln p(y_i, m_i, k_i | \tilde{\mu}_i^{(m)}, \tilde{\tau}_i^{(m)}, \tilde{\mu}_i^{(k)}, \tilde{\tau}_i^{(k)}) \right\rangle_{k_i} + \text{const.} \\ &= \ln N(y_i | m_i, e^{-\langle k_i \rangle}) + \ln N(m_i | \tilde{\mu}_i^{(m)}, \tilde{\tau}_i^{(m)-1}) + \text{const.} \\ &= -\frac{1}{2} \ln(2\pi e^{-\langle k_i \rangle}) - \frac{e^{\langle k_i \rangle} (y_i - m_i)^2}{2} - \frac{1}{2} \ln(2\pi \tilde{\tau}_i^{(m)-1}) - \frac{\tilde{\tau}_i^{(m)} (m_i - \tilde{\mu}_i^{(m)})^2}{2} + \text{const.} \end{aligned} \quad (20)$$

Similarly, the optimal solution for k_i is

$$\begin{aligned} \ln q^*(k_i) &= \left\langle \ln p(y_i, m_i, k_i | \tilde{\mu}_i^{(m)}, \tilde{\tau}_i^{(m)}, \tilde{\mu}_i^{(k)}, \tilde{\tau}_i^{(k)}) \right\rangle_{m_i} + \text{const.} \\ &= \ln N(y_i | \langle m_i \rangle, e^{-k_i}) + \ln N(k_i | \tilde{\mu}_i^{(k)}, \tilde{\tau}_i^{(k)-1}) + \text{const.} \\ &= -\frac{1}{2} \ln(2\pi e^{-k_i}) - \frac{e^{k_i} (y_i - \langle m_i \rangle)^2}{2} - \frac{1}{2} \ln(2\pi \tilde{\tau}_i^{(k)-1}) - \frac{\tilde{\tau}_i^{(k)} (k_i - \tilde{\mu}_i^{(k)})^2}{2} + \text{const.} \end{aligned} \quad (21)$$

Typically, in variational inference one makes use of conjugate prior distributions, which ensure that prior and posterior distributions are of the same type, and furnish straightforward update equations that can be iteratively evaluated (Bishop, 2006; Blei, Kucukelbir, & McAuliffe, 2017). This is not possible here, due to the use of a (non-conjugate) Gaussian prior for the log precision, and the agent thus makes use of gradient ascent, combined with the Laplace approximation, to derive estimates of the posterior mean and variance for each variable (Bishop, 2006; K. Friston, Mattout, Trujillo-Barreto, Ashburner, & Penny, 2007). The use of gradient ascent is particularly attractive here, as it is employed in classical formulations of predictive coding (Friston & Kiebel, 2009). Thus

$$\tilde{\mu}_i^{(m)} = \arg \max_{m_i} (q^*(m_i)) \quad (22)$$

$$\tilde{\tau}_i^{(m)} = -\frac{\partial^2 q^*(m_i)}{\partial m_i^2} \left(\arg \max_{m_i} (q^*(m_i)) \right) \quad (23)$$

$$\tilde{\mu}_i^{(k)} = \arg \max_{k_i} (q^*(k_i)) \quad (24)$$

$$\tilde{\tau}_i^{(k)} = -\frac{\partial^2 q^*(k_i)}{\partial k_i^2} \left(\arg \max_{k_i} (q^*(k_i)) \right) \quad (25)$$

Differentiating $q^*(m_i)$ twice with respect to m_i gives:

$$\frac{\partial q^*(m_i)}{\partial m_i} = e^{\langle k_i \rangle} (y - m_i) - \tilde{\tau}_i^{(m)} (m_i - \tilde{\mu}_i^{(m)}) \quad (26)$$

$$\frac{\partial^2 q^*(m_i)}{\partial m_i^2} = -e^{\langle k_i \rangle} - \tilde{\tau}_i^{(m)}. \quad (27)$$

The first equation is the familiar core of predictive coding, which trades off prediction errors at the first ($y_i - \langle m_i \rangle$) and second ($\langle m_i \rangle - \tilde{\mu}_i^{(m)}$) levels, weighted by their precisions.

Similarly, for the log precision k_i ,

$$\frac{\partial q^*(k_i)}{\partial k_i} = \frac{1}{2} - e^{k_i} \frac{(y_i - \langle m_i \rangle)^2}{2} - \tilde{\tau}_i^{(k)} (k_i - \tilde{\mu}_i^{(k)}), \quad (28)$$

$$\frac{\partial^2 q^*(k_i)}{\partial k_i^2} = -e^{k_i} \frac{(y_i - \langle m_i \rangle)^2}{2} - \tilde{\tau}_i^{(k)}. \quad (29)$$

Gradient ascent is performed using Newton's method of optimisation, in which, for function f at iteration n , variable x is updated such that

$$\Delta x = -\frac{f'(x_n)}{f''(x_n)}. \quad (30)$$

This yields the following coupled equations, which the agent iteratively evaluates:

$$\hat{m}_i = \hat{m}_i - \frac{e^{\hat{k}_i} (y_i - \hat{m}_i) - \tilde{\tau}_i^{(m)} (\hat{m}_i - \tilde{\mu}_i^{(m)})}{-e^{\hat{k}_i} - \tilde{\tau}_i^{(m)}} \quad (31)$$

$$\hat{k}_i = \hat{k}_i - \frac{\frac{1}{2} - e^{\hat{k}_i} \frac{(y_i - \hat{m}_i)^2}{2} - \tilde{\tau}_i^{(k)} (\hat{k}_i - \tilde{\mu}_i^{(k)})}{-e^{\hat{k}_i} \frac{(y_i - \hat{m}_i)^2}{2} - \tilde{\tau}_i^{(k)}} \quad (32)$$

For reasons of computational expedience, in the analyses presented here, we fixed the number of iterations to sixteen, rather than explicitly evaluating convergence.

Design matrix specification

The vectors of trial-by-trial mean and log precision estimates $\hat{\mathbf{m}}$ and $\hat{\mathbf{k}}$ were used to estimate a trial-specific surprise regressor \mathbf{s} as follows:

$$s_i \equiv -\ln \left(N \left(y_i \mid \hat{m}_{i-1}, e^{-\hat{k}_{i-1}} \right) \right), \quad (33)$$

We modelled behaviour using four versions of this agent. In the first (M1: ‘full’) model, both the mean and precision of the distribution were dynamically estimated as described above, using the priors given in Table 1. In the second (M2: ‘precision only’) model, belief updating about the mean was effectively prevented by fixing both the prior mean and variance of the distribution over ε to be 10^{-6} . In the third (M3: ‘mean only’) model by contrast, belief updating about precision was effectively prevented by fixing the prior mean and variance of the distribution over α to be 10^{-6} . In the fourth (M4: ‘fixed’) model, both sorts of belief updating were prevented in a similar fashion. (Note that, where belief updating is prevented, between-subject variability in fixed estimates of the mean or precision is still allowed, since m_0 and k_0 are still fitted as free parameters).

Model	$\sum L$	$\sum L - \sum L_{null}$	Posterior probability	Protected exceedance probability
<i>M1 ('dual estimation')</i>	5358538	49684	0.403	0.441
<i>M2 ('precision only')</i>	5358392	49538	0.427	0.533
<i>M3 ('mean only')</i>	5357825	48971	0.153	0.019
<i>M4 ('fixed')</i>	5357787	48933	0.011	0.004
<i>M0 ('null')</i>	5308854	0	0.007	0.003

Table 1: Model comparison strongly favoured the models in which subjects dynamically updated their beliefs about the precision of the stimulus distribution (M1 and M2), but do not clearly distinguish these two, thus providing no clear evidence about whether subjects also inferred on the mean of the distribution. (Model comparison was performed using model-space averaging, as described in (FitzGerald, Hammerer, Sambrook, & Penny, 2019))

All four GLM-AR models included binary regressors to model tone and target presentation, a regressor encoding stimulus number (to model linear drifts in the pupil responses across a session), and one encoding \mathbf{y} (frequency in log space). (For completeness, we also tested models that included a ‘response’ regressor indicating which stimuli subjects responded to, rather than one indicating target presentation. However, these models provided markedly inferior fits, so we do not discuss them further). In addition, we included a regressor encoding the absolute difference in frequency between each tone and the one immediately preceding it ($|y_i - y_{i-1}|$), to account for simple adaptation effects.

Finally, models M1 and M2 additionally included a regressor encoding $e^{\hat{k}}$ (the trial-by-trial precision estimate). Identical weak (zero mean) priors were set for each regressor, as specified in Table 2.

In addition to these four inference-based models, we included a null (M0) model which included only the AR component of our scheme. This allows us to assess whether any of our models do better than a simple AR process.

2.6. Model-fitting and comparison

Model-fitting was performed using variational Laplace (VL) (Daunizeau, 2017; K. Friston et al., 2007), with prior distributions as specified in Table 2. This is a fast and powerful approach to model-fitting, which furnishes an estimate of the model evidence that is typically more accurate than measures such as the AIC and BIC (W. D. Penny, 2012) but requires that model parameters be treated as Gaussian. We thus transformed our model parameters where necessary, as specified in Table 2. (In addition to

the model-specific parameters described above, VL also estimates a noise log-precision parameter η for each subject, which we report in Table 2). Model comparison was based on the negative variational free energy for each model and subject derived during model fitting.

Parameter	Prior mean (variance)	Group-level posterior (variance)	p -value (corrected)
$\ln(h)$	$\ln(3)$ (2)	1.61 (0.271)	-
$\ln(l)$	$\ln(3)$ (2)	1.59 (0.536)	-
$\ln(d)$	$\ln(0.2)$ (2)	-3.09 (3.390)	-
a	1 (2)	1.00 (<0.001)	-
w_{event}	0 (4)	0.10 (0.114)	1.000
$w_{\text{precision}}$	0 (4)	0.06 (0.027)	1.000
w_{surprise}	0 (4)	0.04 (0.002)	0.004
w_{pitch}	0 (4)	0.00 (<0.001)	1.000
w_{target}	0 (4)	0.54 (0.111)	<0.001
w_{drift}	0 (4)	-0.15 (0.044)	0.027
w_{adapt}	0 (4)	0.00 (<0.001)	1.000
$\ln(\lambda^{(m)})$	$\ln(100)$ (2)	5.72 (13.49)	-
$\ln(\lambda^{(k)})$	$\ln(100)$ (2)	1.31 (2.93)	-
m_0	$\ln(500)$ (2)	9.12 (18.15)	-
$\ln(k_0)$	$\ln(1)$ (2)	5.95 (3.39)	-
η	4 (4)	6.82 (0.475)	-

Table 2: Summary statistics of the prior and group-level posterior distributions over each parameter. (Posterior distributions are based on weighted average single-subject parameter estimates). Parameters were transformed where appropriate to enable use of a Gaussian prior distribution, as required by the VL algorithm ('ln' indicates the natural logarithm). Non-parametric p -values calculated using permutation testing, Bonferroni-corrected for seven comparisons. These provide clear evidence that both surprise and target presentation were consistently associated with pupil dilation, and for a progressive decrease ('drift') in dilation responses over the session.

2.7. Model checking and visualisation

In order to visually compare the accuracy of our model predictions for key task events against observed responses, we preprocessed each subject's pupillometry data by regressing out the AR component of our model, epoching (between -0.2s and 2.5 s) and baseline correcting (for the interval -0.2s to 0s). We then performed identical preprocessing on the predicted time series for each subject, and plotted responses to probe and target tones.

To directly visualise the nature of pupil responses to surprise, and allow comparison with the responses predicted by our modelling, we performed a time-point by time-point regression analysis on the epoched data described above. The regression model contained a constant term, the surprise estimated for each trial from the full model for that subject, and a regressor encoding target trials. This was used to compare predicted and observed responses with each other, as well as with the gamma kernels produced by the GLM-AR modelling.

3. Results

3.1. Behavioural data

Behavioural data were not relevant to our hypothesis, and therefore were only analysed to ensure participants were paying attention to the tones they were exposed to and to check for unexpected effects of block type. Participants had an average hit rate (number of responses to targets over total number of targets) of 0.84 (range = 0.56–0.98) and an average false alarm rate (number of responses to non-targets over total number of non-targets) of 0.008 (range = 0.001 – 0.023). These data suggest that all participants paid attention to the tones, as, despite sometimes missing them, they almost exclusively responded to targets.

We also carried out three paired samples t-tests to investigate whether the type of block (high vs low precision) influenced the hit rate, the false alarm rate and the reaction times. We found no significant difference in hit rate between high ($\mu = 0.85$, $\sigma^2 = 0.015$) and low ($\mu = 0.83$, $\sigma^2 = 0.031$) precision blocks ($t(15) = 0.80$, $P = 0.437$), nor in false alarm rate ($\mu = 0.008$, $\sigma^2 = 0.00004$ for high precision blocks and $\mu = 0.009$, $\sigma^2 = 0.00009$ for low precision blocks, $t(15) = -0.76$, $P = 0.460$). Likewise, we did not find a significant difference in reaction times between high ($\mu = 559$ ms, $\sigma^2 = 8391$) and low ($\mu = 574$ ms, $\sigma^2 = 9701$) precision blocks, though there was some evidence of a trend ($t(15) = -2.01$, $P = 0.063$). In sum, we found no clear evidence for differences in behaviour on the task between blocks.

3.2. Probe tone analysis

We first analysed responses to the standard and deviant probe tones, using a classical model-free analysis. A two-way repeated measures ANOVA revealed no main effect of precision ($F(1,15) = 3.35$, $p = 0.087$), nor of probe type ($F(1,15) = 3.74$, $p = 0.072$). On the other hand, the interaction was significant ($F(1,15) = 24.71$, $p < 0.001$), with the difference in pupil response between deviant probe and standard probe trials being bigger in high precision blocks than in low precision ones (Figure 3). This clearly suggests that subjects learnt about the precision of stimulus distributions, even though these were task irrelevant.

Post-hoc analyses revealed no significant difference between the standard probe trials in high ($\mu = -0.0027$, $\sigma^2 = 0.000045$) and low ($\mu = -0.0009$, $\sigma^2 = 0.000032$) precision blocks ($t(15) = -0.81$, $p = 0.423$), nor between standard ($\mu = -0.0009$, $\sigma^2 = 0.000032$) and deviant ($\mu = -0.0020$, $\sigma^2 = 0.000021$) probe trials in low precision blocks ($t(15) = -0.665$, $p = 0.516$). On the other hand, deviant

probe tones elicited significantly bigger pupil dilation in high ($\mu = 0.0057$, $\sigma^2 = 0.000048$) compared to low ($\mu = -0.0020$, $\sigma^2 = 0.000021$) precision blocks ($t(15) = 4.94$, $p < 0.001$, Bonferroni-corrected). Finally, deviant probes ($\mu = 0.0057$, $\sigma^2 = 0.000048$) resulted in a larger pupil response compared to standard ones ($\mu = -0.0027$, $\sigma^2 = 0.000045$) in high precision blocks, ($t(15) = 3.24$, $p = 0.022$, Bonferroni-corrected). (These results are illustrated in Figure 3). This model-free analysis suggest that larger responses were associated with deviant probes when embedded in a narrower distribution, in keeping with the prediction that these events are more surprising, and consistent with previous work considering electrophysiological and behavioural responses (Garrido et al., 2013).

Our model-free approach, like those in many previous studies of the oddball paradigm, is restricted to consideration of carefully specified probe tones. This is inefficient, because it only considers a small subset of all experimental stimuli, and introduces restrictions that may preclude consideration of more subtle or complex statistical learning effects using pupillometry. Consequently, we performed a complementary model-based analysis, which forms the principle focus of this paper.

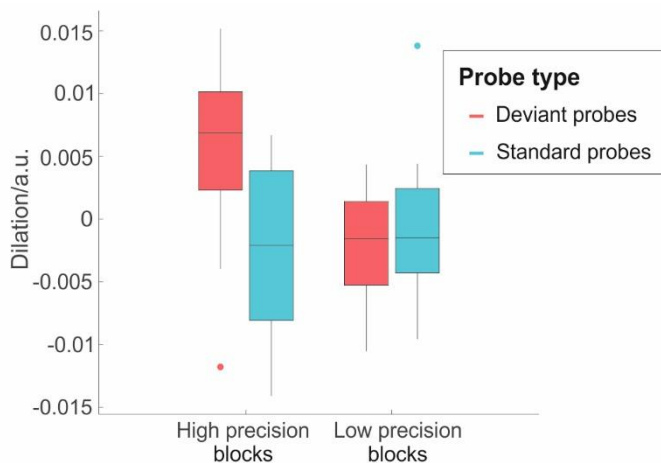


Figure 3: Results of the model-free analysis, illustrating the change in pupil diameter in response to standard (500 Hz) and deviant (2000 Hz) probe tones in the high and low-precision conditions. Greater pupil dilation was observed for deviant compared with standard probes in the high precision condition ($t(15) = 3.24$, $p = 0.022$, Bonferroni-corrected), but not in the low precision condition ($t(15) = -0.665$, $p = 0.516$). This demonstrates that subjects tracked the current precision of the distribution of tones. (Data points between 900ms and 1000ms from stimulus onset were averaged for this analysis. Data are displayed using a Tukey boxplot, with points outside the whiskers ranges additionally plotted).

3.3. GLM-AR modelling

Model comparison was performed using model-space averaging (FitzGerald et al., 2019), a refinement of random-effects Bayesian model selection (Stephan, Penny, Daunizeau, Moran, & Friston, 2009) which automatically (and optimally) mitigates the dilutionary effect of including inferior models in the model-space. This analysis strongly favoured M01 and M02 (the ‘full’ and ‘precision only’ models, Table 1) as compared to the other models, but did not clearly distinguish between them. (These were

assigned posterior probabilities of 0.403 and 0.427, and protected exceedance probabilities (FitzGerald et al., 2019; Rigoux, Stephan, Friston, & Daunizeau, 2014) of 0.441 and 0.533). This provides clear evidence that subjects tracked the precision of the tone distribution, even though it was task-irrelevant, but does not settle the issue of whether they also tracked the mean of the distribution. (This is perhaps unsurprising, given that the mean actually remained constant across the experiment, and could usefully be investigated in future work). We thus selected M02 for use when performing analyses of model performance as described below, on the basis that this seemed the ‘conservative’ option.

One possible alternative explanation for our precision-tracking results is that responses might reflect a similarity effect of recent stimuli, which might be expected to differ between conditions. (In the high precision condition recent stimuli tend to be more similar in frequency to the current stimulus than in the low precision condition). We control for this in our main analysis through the use of an adaptation regressor encoding the absolute difference between the current and previous stimulus, but, as pointed out by a reviewer, this might not be sufficient if the similarity effect involved multiple recent stimuli. To rule this out, we carried out an extra check analysis in which we included separate regressors encoding the absolute difference between the current stimulus and each of the seven preceding ones. We used this augmented approach to compare the ‘full’ and ‘mean only’ models, to see if there was evidence in favour of precision tracking even when this fuller stimulus history was accounted for. Reassuringly, this also provided strong evidence in favour of the full model, which had an exceedance probability of 0.994.

Quality of model fits, as estimated using simple percentage variance explained, was excellent ($\mu = 0.999$, range: 0.997 – 1.000). However, this includes the effects of the AR process, which is not of interest here, so it is also useful to assess how much variance is explained by the GLM itself. To assess this, we regressed out the predictions of the AR component from each subject’s time series, and then calculated percentage variance explained solely by the GLM. This also showed a good fit with the data ($\mu = 0.848$, range: 0.154 – 0.990), as illustrated by Figures 4 and 5. Inspection of the convolution kernels derived from our modelling suggested that these provided plausible pupil dilation responses both when compared with existing literature (Denison et al., 2019; Hong et al., 2014; Knapen et al., 2016; Korn & Bach, 2016; Murphy et al., 2011), and when compared with the waveform generated when regressing estimated beliefs about surprise onto epoched data (Figure 5). This suggests that our GLM-AR modelling approach was appropriate for analysing these data, and supports its use in future studies.

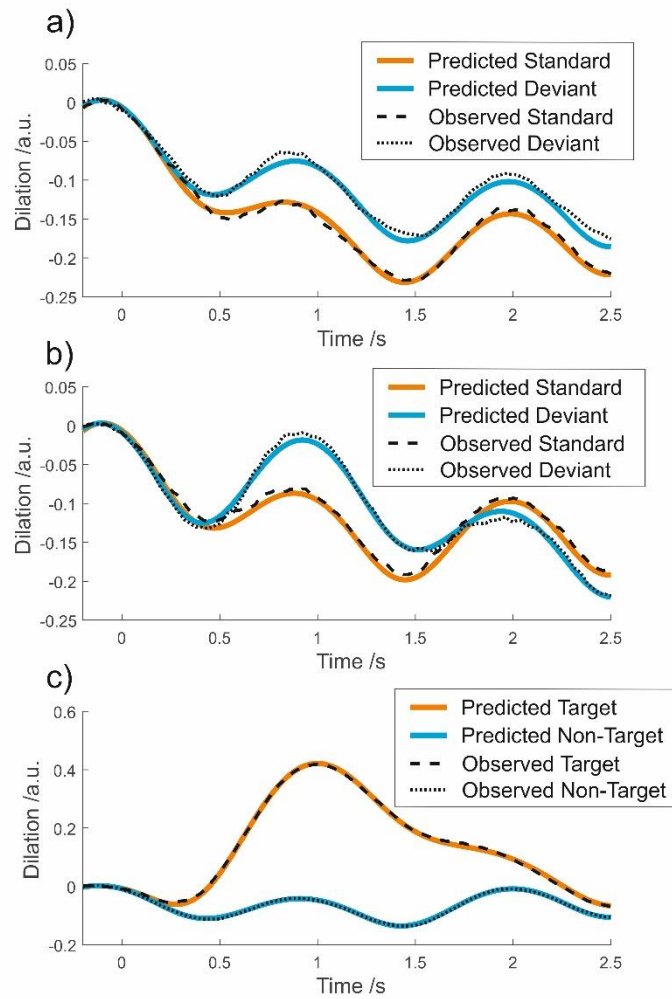


Figure 4: Illustration of the accuracy of GLM-derived predictions for a single representative subject (subject two, percentage variance explained = 0.87, similar to the group mean). Each plot shows epoched, baseline corrected and averaged waveforms for the predicted (blue/orange) and observed (black) responses to key task conditions. (The predictions of the AR component of the model were regressed out prior to epoching, and these thus solely reflect how well the GLM predicts the data). The top plot (a) illustrates probe tones in the low precision condition, the middle plot (b) illustrates probe tones in the high precision condition, and the bottom plot (c) illustrates target and non-target tones (across both conditions). For all waveforms there is a close correspondence between predicted and observed data, reflecting the accuracy of the model fits.

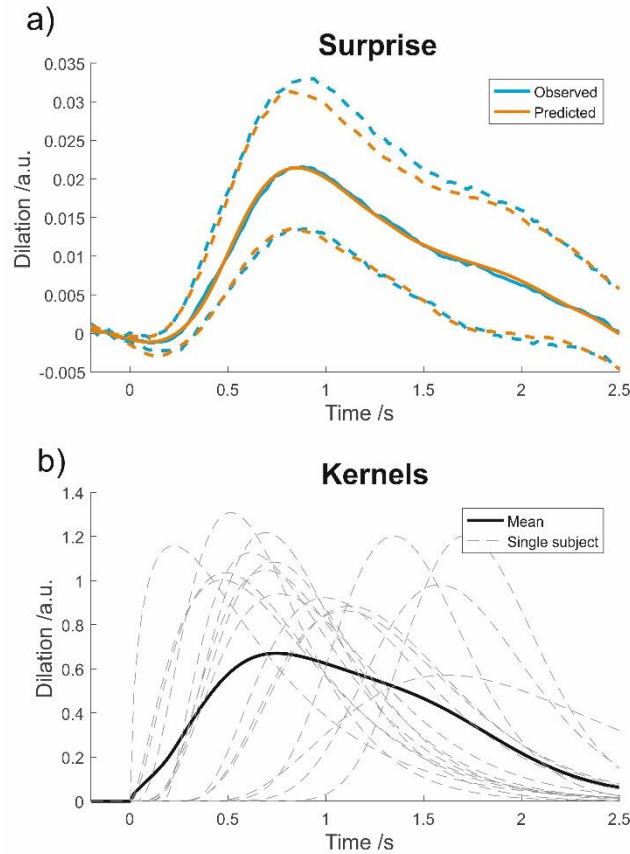


Figure 5a: Surprise waveform estimated from our regression analysis (see Methods for more details). Observed dilation responses to surprise (blue) peak at roughly one second and then return to baseline. Predicted responses to surprise derived from our GLM-AR modelling (orange) show a close correspondence to observed responses. (Dashed lines indicate bootstrapped 95% confidence intervals). 5b: Gamma kernels modelling pupil dilation derived from the GLM-AR model. (Single subject responses in grey, and the mean in black). These strongly resemble both the surprise waveform derived in our regression analysis (4a) and averaged responses from tasks using slower designs (for example (Hong et al., 2014)).

In addition, we explored how closely the predictions made by the fitted models tracked the true precision of the stimulus distributions (Figure 6). Analysis using Spearman's rank correlation coefficient showed a strong positive correlation between trial-by-trial precision estimated and the true task contingencies ($\bar{\rho} = 0.77$, $\sigma^2 = 0.01$), suggesting that our cognitive model performed adequately on the task.




Figure 6: Illustration of how trial-by-trial estimates of precision (blue) track the true precision of the distribution used to generate the non-probe tones (orange). Data is plotted from the first 1200 trials in a single representative subject (subject two). The tendency to underestimate precision in the high-precision blocks most likely reflects the distorted (and lower precision) probability distributions induced by the use of probe tones.

To test for group-level effects of the factors in the GLM, we calculated an average of the maximum a posteriori (MAP) parameter estimates for each subject, weighted by the posterior probability assigned to each model in that subject during the model comparison. Inference was then performed using a permutation test in which we flipped the signs of the parameters in a randomly selected subset of subjects 100,000 times and used the resultant surrogate data to provide surrogate two-tailed p -values for each variable. (These were corrected for seven comparisons using a Bonferroni correction). Clear evidence of positive dilation responses to surprise ($\bar{w} = 0.040$, $p = 0.004$) and target presentation ($\bar{w} = 0.538$, $p < 0.001$) were found, as well as a negative effect of trial order ($\bar{w} = -0.154$, $p = 0.026$), suggesting a progressive decrease in the size of dilation responses across the course of the experiment (Figure 7). No statistically significant effects were observed for precision itself, log frequency, tone presentation, or the frequency separation between successive tones. (See Table 2 for full results). Additionally, we tested for between-subject correlations in the regression coefficients, using a partial correlation approach to control for non-specific differences in coefficient magnitude. (These might be caused, for example, by quality of model fit, or the subject-specific shape of the gamma kernel). This showed no evidence for statistically significant correlations between coefficients, and since we have no clear hypotheses about such relationships, we do not discuss them further.

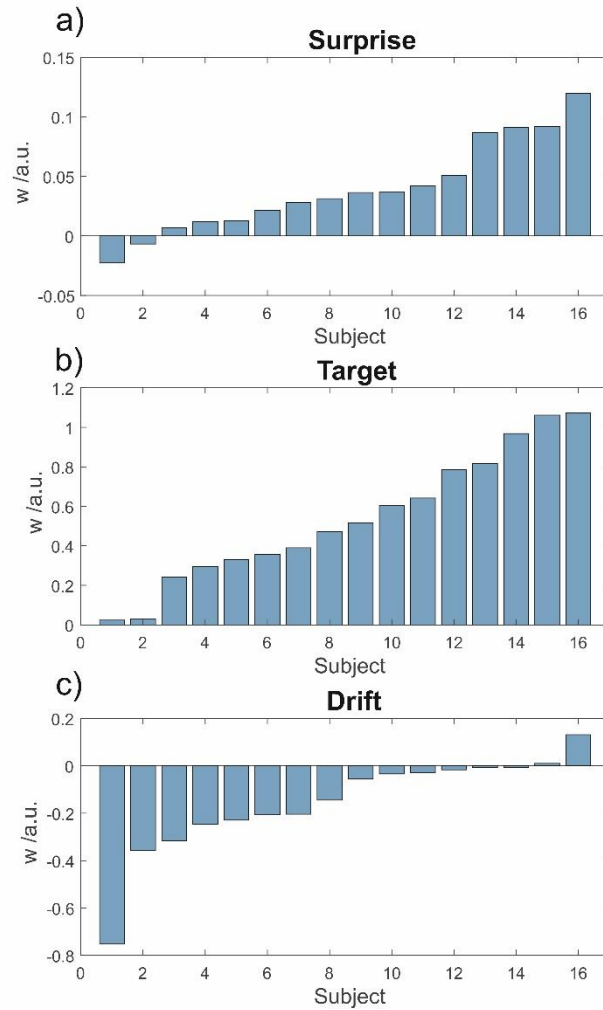


Figure 7: Ordered parameter estimates of single subject regression weights for the Surprise regressor (w_{surprise} , 7a), Target regressor (w_{target} , 7b), and linear drift (w_{drift} , 7c) derived using weighted averaging (see Methods for further details). Positive pupil dilation responses to both surprise and target presentation were highly consistent across subjects, as was a progressive decrease in the size of responses over time.

In line with previous work (Garrido et al., 2013), the stimulus distributions that we used in this experiment are distorted to introduce probe tones that can be compared across conditions (Figure 2). This introduces the possibility that these tones are treated differently by subjects, and might be responsible for driving our results. To rule this out, we repeated the model comparison described above, using models that ignored probe trials. Reassuringly, these results were very similar, with M1 and M2 assigned posterior probabilities of 0.382 and 0.443 respectively, and protected exceedance probabilities of 0.370 and 0.604. This suggests that our key results were not driven by responses to the probe tones.

4. Discussion

The results obtained in this study provide clear evidence that pupil dilation reflects automatic and dynamically updated beliefs about the precision of stimulus distributions, in keeping with theories of probabilistic cognition (Aitchison & Lengyel, 2016; Friston, 2010; Ma et al., 2006; Tenenbaum, Griffiths, & Kemp, 2006). This extends previous work showing evidence for an effect of the precision of stimulus distributions on reaction times and MEG responses (Garrido et al., 2013), and suggests that pupillometry can be a useful tool for examining statistical learning about higher order properties of stimulus distributions (Alamia et al., 2019), something we will consider further in future work. In addition, we use these data to demonstrate the potential for analysing pupillometry data using a GLM-AR approach, which allowed highly accurate prediction of observed responses in our data (Figures 4, 5).

Our study complements an existing body of work using pupillometry to explore learning and related processes (Denison et al., 2019; Ebitz & Moore, 2018; Mathôt, 2018). Perhaps the simplest manifestation of this is in the literature on pupil dilation in response to perceptual oddballs (Friedman et al., 1973; Hong et al., 2014; Korn & Bach, 2016; Liao et al., 2016; Murphy et al., 2011; Qiyuan et al., 1985; Steinhauer & Zubin, 1982), but a similar approach has been adopted to explore response during gambling and learning tasks (Hämmerer et al., 2019; Lavín et al., 2014; Preuschoff et al., 2011), change-point detection (Nassar et al., 2012), the role of risk and learning about transition probabilities between discrete states (Alamia et al., 2019), and responses to volatility (Browning et al., 2015; Vincent et al., 2019), as well as surprise in other contexts (e.g. (Kloosterman et al., 2015; Knapen et al., 2016; O'Reilly et al., 2013)). As such, the principle contribution of our findings is to provide new information about the sort of cognitive processes that are reflected in pupil dilation responses, and contribute to the growing literature linking them specifically to statistical learning (Alamia et al., 2019).

A key limitation of many previous studies using oddball paradigms is the conflation of deviant and target stimuli (Hong et al., 2014; Liao et al., 2016; Murphy et al., 2014; Murphy, Robertson, Balsters, & O'Connell, 2011; J Rajkowski et al., 1994; Janusz Rajkowski et al., 2004; Steinhauer & Zubin, 1982 (though see Wetzell, Buttellmann, Schieler, & Widmann, 2016 for studies which avoid this). This conflation makes it difficult to attribute pupil/LC effects unequivocally to surprise rather than other task-related processes. We obviated this by making pitch irrelevant to the task, so that evoked pupil dilation could be directly associated with stimulus probability (and, therefore, surprise). Importantly, having pitch be task-irrelevant allowed us to explore automatic, and possibly implicit, learning processes. This provides a complement to paradigms in which learning is directly relevant for behaviour, and indexes what is likely to be an important form of learning for behaviour in ecological settings.

It should be noted that a recent review (Zenon, 2019), in an attempt to give a unified explanation to the pupil effects of a wide range of cognitive processes (e.g. mental effort, attention, exploration/exploitation trade-off, decision making, surprise), related pupil dilation to information gain. This is formalised as the Kullback–Leibler divergence between prior and posterior distributions (often called 'Bayesian surprise' (Baldi & Itti, 2010; Philipp Schwartenbeck et al., 2016), as opposed to information-theoretic ('Shannon') surprise (negative log probability). Bayesian surprise has the advantage that it quantifies the meaningful information present in a stimulus, as opposed to simply how unexpected it is, and it is thus a plausible candidate to play a role here. However, our paradigm is not designed to separate these two quantities, and thus evaluating them as competing explanations for pupil dilation responses goes beyond the scope of our study. Similarly, given the close relationship

between surprise and dynamic beliefs about volatility (Silvetti, Seurinck, van Bochove, & Verguts, 2013), it is conceivable that the dilation responses we observe more directly index beliefs about volatility, but this falls outside the scope of our study to test.

A diverse set of evidence points to a tight link between pupil diameter and noradrenergic activity in locus coeruleus (LC). This was first observed in monkeys, with an electrophysiological study showing that pupil diameter tracked LC tonic activity (J Rajkowski, 1993). Pharmacological evidence confirmed this finding in humans, with LC-suppressing drugs decreasing pupil diameter and LC-stimulating drugs enhancing it (Hou, Freeman, Langley, Szabadi, & Bradshaw, 2005; Phillips, Szabadi, & Bradshaw, 2000). In addition, theoretical work (Dayan & Yu, 2006) has linked LC activity with surprise, as electrophysiological data on monkeys seem to suggest (J Rajkowski et al., 1994; Janusz Rajkowski et al., 2004). Murphy and colleagues carried out an experiment linking everything together, showing how LC BOLD activity correlates with pupil diameter in resting state and how they respond similarly to deviant stimuli in an oddball task (Murphy et al., 2014), supporting the already popular idea (e.g. Lavin et al., 2014; Preuschoff et al., 2011) that surprise-related pupil dilation occurs as an effect of phasic noradrenergic activity. Assuming that this link holds here, our study thus makes the novel contribution that noradrenergic function (and thus the cognitive processes it subserves (Dayan & Yu, 2006)) are sensitive to dynamically updated beliefs about stimulus precision, and may even play a role in this process.

Use of auditory oddball paradigms to index automatic statistical learning has considerable practical attractions, not least the fact that it requires minimal subject compliance (Boly et al., 2011). Combining this with pupillometric data collection is attractive, as such data is relatively simple and cheap to collect, particularly when compared with neuroimaging modalities such as MEG and fMRI. The general approach here thus has potential for exploring cognitive changes related to statistical learning in patient groups (Browning et al., 2015), as well as during healthy ageing (Hämmerer et al., 2019).

The GLM-AR approach that we adopt here, though originally motivated by the strong similarities between pupillometry and fMRI data (W. Penny et al., 2003), relates closely to a number of existing approaches. The use of a convolution kernel for analysing pupillometric data was first proposed by Hoeks and Levelt (Hoeks & Levelt, 1993), and similar approaches have been adopted by various authors subsequently (e.g. Wierda et al., 2012; de Gee et al., 2014; Knapen et al., 2016; Korn and Bach, 2016; Denison et al., 2019; Vincent et al., 2019). (We note though that many of these studies make use of a canonical pupil response, which is likely to be suboptimal, given the evidence for individual variability (Denison et al., 2019)). However, such approaches do not account for the strong slow fluctuations in pupil diameter (Zénon, 2017). Typically, the effects of these fluctuations are mitigated via epoching, baseline-correction, and averaging, but this requires many repetitions of a particular trial type, which makes it difficult to capture phenomena such as learning (as we do here). We address this issue through use of an AR model, as has been proposed recently by Zenon (Alamia et al., 2019; Zénon, 2017). However, to our knowledge the GLM-AR approach we propose is the first to combine the advantages of both individually tailored convolution kernels and AR modelling. A rigorous assessment of the significance of this is beyond the scope of this paper, but we will explore it in future work.

A further, (and critical for our purposes), aspect of the GLM-AR approach that we use is that it allows us to fit the parameters of behavioural models to pupillometric responses an approach more typically confined to analysis of behavioural data (for example (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; O'Doherty et al., 2007; P Schwartenbeck, FitzGerald, Mathys, Dolan, & Friston, 2015; Smittenaar, FitzGerald, Romei, Wright, & Dolan, 2013)). This 'doubly model-based' aspect is

important, as it allows us to use model comparison to establish which behavioural models best account for pupillometric data. (For example, one might use it to compare different models of learning). In addition, it permits, in principle, the characterisation of between-subject variability in processes such as learning and inference, which may be of particular interest for understanding pathology (Browning et al., 2015; Huys, Maia, & Frank, 2016; Krystal et al., 2017; Montague et al., 2012). However, the extent to which applying the GLM-AR approach to pupillometric data in practice permits such inferences is unclear, and future work will be necessary to establish this.

In sum, at a cognitive level, our work represents a contribution both to understanding task-irrelevant human statistical learning processes, and to characterising the computational mechanisms underlying pupil dilation responses to surprising stimuli. In addition, we believe that the GLM-AR approach that we propose has considerable potential for increasing the accuracy and flexibility of pupillometry data analysis, something we will explore in future work.

Declarations of interest: None.

Acknowledgments

THBF is supported by a European Research Council (ERC) Starting Grant under the Horizon 2020 program (Grant Agreement 804701). FS is supported by a studentship from the School of Psychology, UEA. We thank Richard Carey-Knight for technical support with pupillometry data collection. Pupillometry data are available on request from the authors. The code for the GLM-AR model is available at <https://github.com/thbfitz/pupil>.

References

- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013). The computational anatomy of psychosis. *Frontiers in Psychiatry, 4*, 47. <https://doi.org/10.3389/fpsyt.2013.00047>
- Aitchison, L., & Lengyel, M. (2016). The Hamiltonian Brain: Efficient Probabilistic Inference with Excitatory-Inhibitory Neural Circuit Dynamics. *PLoS Computational Biology, 12*(12), 1–24. <https://doi.org/10.1371/journal.pcbi.1005186>
- Alamia, A., VanRullen, R., Pasqualotto, E., Mouraux, A., & Zenon, A. (2019). Pupil-Linked Arousal Responds to Unconscious Surprisal. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 39*(27), 5369–5376. <https://doi.org/10.1523/JNEUROSCI.3010-18.2019>
- Baldi, P., & Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks : The Official Journal of the International Neural Network Society, 23*(5), 649–666. <https://doi.org/10.1016/j.neunet.2009.12.007>
- Beal, M. J. (2003). *Variational algorithms for approximate Bayesian inference*. University of London.
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience, 10*(9), 1214–1221. <https://doi.org/10.1038/nn1954>
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association, 0*–0.

<https://doi.org/10.1080/01621459.2017.1285773>

- Boly, M., Garrido, M. I., Gosseries, O., Bruno, M.-A., Boveroux, P., Schnakers, C., ... Friston, K. (2011). Preserved Feedforward But Impaired Top-Down Processes in the Vegetative State. *Science*, 332(6031), 858–862. <https://doi.org/10.1126/science.1202043>
- Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, 18(4), 590–596. <https://doi.org/10.1038/nn.3961>
- Damsa, A., & van Rijn, H. (2017). Pupillary response indexes the metrical hierarchy of unattended rhythmic violations. *Brain and Cognition*, 111, 95–103. <https://doi.org/10.1016/j.bandc.2016.10.004>
- Daunizeau, J. (2017). The variational Laplace approach to approximate Bayesian inference. *ArXiv*, 1703.02089.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron*, 69(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>
- Dayan, P., & Yu, A. J. (2006). Phasic norepinephrine: a neural interrupt signal for unexpected events. *Network: Computation in Neural Systems*, 17(4), 335–350.
- De Berker, A. O., Rutledge, R. B., Mathys, C., Marshall, L., Cross, G. F., Dolan, R. J., & Bestmann, S. (2016). Computations of uncertainty mediate acute stress responses in humans. *Nature Communications*, 7, 10996.
- de Gee, J. W., Knapen, T., & Donner, T. H. (2014). Decision-related pupil dilation reflects upcoming choice and individual bias. *Proceedings of the National Academy of Sciences of the United States of America*, 111(5), E618-25. <https://doi.org/10.1073/pnas.1317557111>
- Denison, R. N., Parker, J. A., & Carrasco, M. (2019). Modeling pupil responses to rapid sequential events. *BioRxiv*, 655902. <https://doi.org/10.1101/655902>
- Ebitz, R. B., & Moore, T. (2018). Both a Gauge and a Filter: Cognitive Modulations of Pupil Size. *Frontiers in Neurology*, 9, 1190. <https://doi.org/10.3389/fneur.2018.01190>
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, Vol. 14, pp. 119–130. <https://doi.org/10.1016/j.tics.2010.01.003>
- FitzGerald, T. H., Hammerer, D., Sambrook, T. D., & Penny, W. D. (2019). Bayesian inference over model-spaces increases the accuracy of model comparison and allows formal testing of hypotheses about model distributions in experimental populations. *ArXiv*, 1901.01916.
- Fletcher, P., & Frith, C. (2008). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews Neuroscience*, 10(1), 48–58.
- Friedman, D., Hakerem, G., Sutton, S., & Fleiss, J. L. (1973). Effect of stimulus uncertainty on the pupillary dilation response and the vertex evoked potential. *Electroencephalography and Clinical Neurophysiology*, 34(5), 475–484.
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4(11), e1000211. <https://doi.org/10.1371/journal.pcbi.1000211>
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>

- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory. *Neural Computation*, *29*(1), 1–49. https://doi.org/10.1162/NECO_a_00912
- Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: the brain as a phantastic organ. *The Lancet. Psychiatry*, *1*(2), 148–158. [https://doi.org/10.1016/S2215-0366\(14\)70275-5](https://doi.org/10.1016/S2215-0366(14)70275-5)
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1521), 1211–1221. <https://doi.org/10.1098/rstb.2008.0300>
- Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., & Penny, W. (2007). Variational free energy and the Laplace approximation. *NeuroImage*, *34*(1), 220–234. <https://doi.org/10.1016/j.neuroimage.2006.08.035>
- Garrido, M. I., Sahani, M., & Dolan, R. J. (2013). Outlier Responses Reflect Sensitivity to Statistical Structure in the Human Brain. *PLoS Computational Biology*, *9*(3). <https://doi.org/10.1371/journal.pcbi.1002999>
- Hämmerer, D., Schwartenbeck, P., Gallagher, M., FitzGerald, T. H. B., Düzel, E., & Dolan, R. J. (2019). Older adults fail to form stable task representations during model-based reversal inference. *Neurobiology of Aging*, *74*, 90–100. <https://doi.org/10.1016/j.neurobiolaging.2018.10.009>
- Hoeks, B., & Levelt, W. J. M. (1993). Pupillary dilation as a measure of attention: a quantitative system analysis. *Behavior Research Methods, Instruments, & Computers*, *25*(1), 16–26. <https://doi.org/10.3758/BF03204445>
- Hong, L., Walz, J. M., & Sajda, P. (2014). Your eyes give you away: Prestimulus changes in pupil diameter correlate with poststimulus task-related EEG dynamics. *PLoS ONE*, *9*(3). <https://doi.org/10.1371/journal.pone.0091321>
- Hou, R. H., Freeman, C., Langley, R. W., Szabadi, E., & Bradshaw, C. M. (2005). Does modafinil activate the locus coeruleus in man? Comparison of modafinil and clonidine on arousal and autonomic functions in human volunteers. *Psychopharmacology*, *181*(3), 537–549.
- Huys, Q. J. M., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, *19*(3), 404–413. <https://doi.org/10.1038/nn.4238>
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*(2), B35–B42.
- Kloosterman, N. A., Meindertsma, T., van Loon, A. M., Lamme, V. A. F., Bonneh, Y. S., & Donner, T. H. (2015). Pupil size tracks perceptual content and surprise. *European Journal of Neuroscience*, *41*(8), 1068–1078.
- Knapen, T., de Gee, J. W., Brascamp, J., Nuiten, S., Hoppenbrouwers, S., & Theeuwes, J. (2016). Cognitive and Ocular Factors Jointly Determine Pupil Responses under Equiluminance. *PLoS One*, *11*(5), e0155574. <https://doi.org/10.1371/journal.pone.0155574>
- Korn, C. W., & Bach, D. R. (2016). A solid frame for the window on cognition: Modeling eventrelated pupil responses. *Journal of Vision*, *16*(3), 1–16. <https://doi.org/10.1167/16.3.28>
- Krystal, J. H., Murray, J. D., Chekroud, A. M., Corlett, P. R., Yang, G., Wang, X.-J., & Anticevic, A. (2017). Computational Psychiatry and the Challenge of Schizophrenia. *Schizophrenia Bulletin*, *43*(3), 473–475. <https://doi.org/10.1093/schbul/sbx025>
- Lavín, C., Martín, R. S., & Jubal, E. R. (2014). Pupil dilation signals uncertainty and surprise in a

- learning gambling task. *Frontiers in Behavioral Neuroscience*, 7(JAN), 1–8.
<https://doi.org/10.3389/fnbeh.2013.00218>
- Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00302>
- Liao, H.-I., Yoneya, M., Kidani, S., Kashino, M., & Furukawa, S. (2016). Human Pupillary Dilation Response to Deviant Auditory Stimuli: Effects of Stimulus Properties and Voluntary Attention. *Frontiers in Neuroscience*, 10, 43. <https://doi.org/10.3389/fnins.2016.00043>
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432–1438. <https://doi.org/10.1038/nn1790>
- Mathôt, S. (2018). Pupillometry: Psychology, Physiology, and Function. *Journal of Cognition*, 1(1). <https://doi.org/10.5334/joc.18>
- Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience*, 5, 39. <https://doi.org/10.3389/fnhum.2011.00039>
- Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational psychiatry. *Trends in Cognitive Sciences*, 16(1), 72–80. <https://doi.org/10.1016/j.tics.2011.11.018>
- Murphy, P. R., O’Connell, R. G., O’Sullivan, M., Robertson, I. H., & Balsters, J. H. (2014). Pupil diameter covaries with BOLD activity in human locus coeruleus. *Human Brain Mapping*, 35(8), 4140–4154.
- Murphy, P. R., Robertson, I. H., Balsters, J. H., & O’Connell, R. G. (2011). Pupillometry and P3 index the locus coeruleus-noradrenergic arousal function in humans. *Psychophysiology*, 48(11), 1532–1543. <https://doi.org/10.1111/j.1469-8986.2011.01226.x>
- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, 15(7), 1040.
- O’Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-Based fMRI and Its Application to Reward Learning and Decision Making. *Annals of the New York Academy of Sciences*, 1104(1), 35–53. <https://doi.org/10.1196/annals.1390.022>
- O’Reilly, J. X., Schuffelgen, U., Cuell, S. F., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 110(38), E3660–E3669. <https://doi.org/10.1073/pnas.1305373110>
- Penny, W. D. (2012). Comparing dynamic causal models using AIC, BIC and free energy. *NeuroImage*, 59(1), 319–330. <https://doi.org/10.1016/j.neuroimage.2011.07.039>
- Penny, W., Kiebel, S., & Friston, K. (2003). Variational Bayesian inference for fMRI time series. *NeuroImage*, 19(3), 727–741. [https://doi.org/10.1016/S1053-8119\(03\)00071-5](https://doi.org/10.1016/S1053-8119(03)00071-5)
- Phillips, M. A., Szabadi, E., & Bradshaw, C. M. (2000). Comparison of the effects of clonidine and yohimbine on spontaneous pupillary fluctuations in healthy human volunteers. *Psychopharmacology*, 150(1), 85–89.
- Preuschoff, K., ’t Hart, B. M., & Einhäuser, W. (2011). Pupil dilation signals surprise: Evidence for noradrenaline’s role in decision making. *Frontiers in Neuroscience*, 5(SEP), 1–12. <https://doi.org/10.3389/fnins.2011.00115>

- Qiyuan, J., Richer, F., Wagoner, B. L., & Beatty, J. (1985). The pupil and stimulus probability. *Psychophysiology*, 22(5), 530–534.
- Raisig, S., Welke, T., Hagendorf, H., & van der Meer, E. (2010). I spy with my little eye: Detection of temporal violations in event sequences and the pupillary response. *International Journal of Psychophysiology*, 76(1), 1–8.
- Rajkowski, J. (1993). Correlations between locus coeruleus (LC) neural activity, pupil diameter and behavior in monkey support a role of LC in attention. *Soc. Neurosc., Abstract, Washington, DC, 1993*.
- Rajkowski, J., Kubiak, P., & Aston-Jones, G. (1994). Locus coeruleus activity in monkey: phasic and tonic changes are associated with altered vigilance. *Brain Research Bulletin*, 35(5–6), 607–616.
- Rajkowski, Janusz, Majczynski, H., Clayton, E., & Aston-Jones, G. (2004). Activation of monkey locus coeruleus neurons varies with difficulty and performance in a target detection task. *Journal of Neurophysiology*, 92(1), 361–371.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.
<https://doi.org/10.1038/4580>
- Reinhard, G., & Lachnit, H. (2002). The effect of stimulus probability on pupillary response as an indicator of cognitive processing in human learning and categorization. *Biological Psychology*, 60(2–3), 199–215.
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies - Revisited. *NeuroImage*, 84, 971–985.
<https://doi.org/10.1016/j.neuroimage.2013.08.065>
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27–52.
- Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., & Friston, K. (2015). The Dopaminergic Midbrain Encodes the Expected Certainty about Desired Outcomes. *Cerebral Cortex*, 25(10), 3434–3445. <https://doi.org/10.1093/cercor/bhu159>
- Schwartenbeck, Philipp, FitzGerald, T. H. B., & Dolan, R. J. (2016). Neural signals encoding shifts in beliefs. *NeuroImage*, 125, 578–586. <https://doi.org/10.1016/j.neuroimage.2015.10.067>
- Shipp, S. (2016). Neural elements for predictive coding. *Frontiers in Psychology*, Vol. 7.
<https://doi.org/10.3389/fpsyg.2016.01792>
- Silvetti, M., Seurinck, R., van Bochove, M., & Verguts, T. (2013). The influence of the noradrenergic system on optimal control of neural plasticity. *Frontiers in Behavioral Neuroscience*, 7, 160.
- Smittenaar, P., FitzGerald, T. H. B., Romei, V., Wright, N., & Dolan, R. J. (2013). Disruption of dorsolateral prefrontal cortex decreases model-based in favour of model-free control in humans. *Neuron*, 80(4), 914–919.
- Spratling, M. W. (2017). A review of predictive coding algorithms. *Brain and Cognition*, 112, 92–97.
<https://doi.org/10.1016/J.BANDC.2015.11.003>
- Steinhauer, S., & Zubin, J. (1982). Vulnerability to schizophrenia: Information processing in the pupil and event-related potential. In *Biological markers in psychiatry and neurology* (pp. 371–385).

Elsevier.

- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*(4), 1004–1017. <https://doi.org/10.1016/j.neuroimage.2009.03.025>
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318. <https://doi.org/10.1016/j.tics.2006.05.009>
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science (New York, N.Y.)*, *331*(6022), 1279–1285. <https://doi.org/10.1126/science.1192788>
- Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(33), 11177–11187. <https://doi.org/10.1523/JNEUROSCI.0858-10.2010>
- Vincent, P., Parr, T., Benrimoh, D., & Friston, K. J. (2019). With an eye on uncertainty: Modelling pupillary responses to environmental volatility. *PLOS Computational Biology*, *15*(7), e1007126. <https://doi.org/10.1371/journal.pcbi.1007126>
- Wetzel, N., Buttelmann, D., Schieler, A., & Widmann, A. (2016). Infant and adult pupil dilation in response to unexpected sounds. *Developmental Psychobiology*, *58*(3), 382–392.
- Wierda, S. M., van Rijn, H., Taatgen, N. A., & Martens, S. (2012). Pupil dilation deconvolution reveals the dynamics of attention at high temporal resolution. *Proceedings of the National Academy of Sciences*, *109*(22), 8456–8460. <https://doi.org/10.1073/pnas.1201858109>
- Zenon, A. (2019). Eye pupil signals information gain. *BioRxiv*, 693838.
- Zénon, A. (2017). Time-domain analysis for extracting fast-paced pupil responses. *Scientific Reports*, *7*(1), 41484. <https://doi.org/10.1038/srep41484>